

Behavioral Mechanism Design: Optimal Crowdsourcing Contracts and Prospect Theory

DAVID EASLEY, Cornell University
ARPITA GHOSH, Cornell University

Incentive design is more likely to elicit desired outcomes when it is derived based on accurate models of agent behavior. A substantial literature in behavioral economics, however, demonstrates that individuals systematically and consistently deviate from the standard economic model—expected utility theory—for decision-making under uncertainty, which is at the core of the equilibrium analysis necessary to facilitate mechanism design. Can these behavioral biases—as modeled by *prospect theory* [Kahneman and Tversky 1979]—in agents’ decision-making make a difference to the optimal design of incentives in these environments? In this paper, we explore this question in the context of markets for online labor and crowdsourcing where workers make strategic choices about whether to undertake a task, but do not strategize over quality conditional on participation. We ask what kind of incentive scheme—amongst a broad class of contracts, including those observed on major crowdsourcing platforms such as fixed prices or base payments with bonuses (as on MTurk or oDesk), or open-entry contests (as on platforms like Kaggle or Topcoder)—a principal might want to employ, and how the answer to this question depends on whether workers behave according to expected utility or prospect theory preferences.

We first show that with expected utility agents, the optimal contract—for any increasing objective of the principal—always takes the form of an *output-independent* fixed payment to some optimally chosen number of agents. In contrast, when agents behave according to prospect theory preferences, we show that a winner-take-all *contest* can dominate the fixed-payment contract, for large enough total payments, under a certain condition on the preference functions; we show that this condition is satisfied for the parameters given by the literature on econometric estimation of the prospect theory model [Tversky and Kahneman 1992; Bruhin et al. 2010]. Since these estimates are based on fitting the prospect theory model to extensive experimental data, this result provides a strong affirmative answer to our question for ‘real’ population preferences: a principal might indeed choose a fundamentally different kind of mechanism—an output-contingent contest versus a ‘safe output-independent scheme—and do better as a result, if he accounts for deviations from the standard economic models of decision-making that are typically used in theoretical design.

Categories and Subject Descriptors: J.4 [Computer Applications]: Social & Behavioral Sciences.

Additional Key Words and Phrases: Contracts; Contests; Crowdsourcing; Mechanism design; Behavioral; Game theory.

1. INTRODUCTION

The vast range of systems with outcomes that depend on the choices made by economic agents has led to a rich and large literature on mechanism design, which regards designing incentives so that agents make choices resulting in ‘good’ outcomes. Incentives are more likely to elicit desired outcomes when they are derived based on accurate models of agent behavior. A growing literature, however, suggests that people do not quite behave like the standard economic agents in the mechanism design literature. Can such differences have significant consequences for the optimal design of incentive mechanisms?

Decision making under uncertainty: Prospect theory. A particular instance of such a difference involves behavioral biases in decision-making under uncertainty. Many, if not most, environments to which the mechanism design literature has been applied involve risky choice—agents who must

Authors addresses: D. Easley, Departments of Economics and Information Science, Cornell University, email: dae3@cornell.edu; A. Ghosh, Department of Information Science, Cornell University, email: ag865@cornell.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC’15, June 15–19, 2015, Portland, OR, USA.

ACM 978-1-4503-3410-5/15/06 ...\$15.00.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

<http://dx.doi.org/10.1145/http://dx.doi.org/10.1145/http://dx.doi.org/10.1145/2764468.2764513>

make decisions between a set of choices, each of which yields an associated payoff with some probability. The standard economic model for choice under uncertainty is *expected utility theory* [von Neumann and Morgenstern 1944; Savage 1954], whereby agents make choices by comparing the expected utility of outcomes resulting from each possible choice. A substantial body of work in behavioral economics, however, demonstrates that individuals display systematic deviations from the expected utility model, including overweighting low-probability events and under-weighting high-probability ones, as well as displaying a greater disutility for losses compared to utilities from a gain of the same amount.

Prospect theory, introduced in Kahneman and Tversky [1979] and further refined to cumulative prospect theory in Tversky and Kahneman [1992], is a descriptive model that describes much empirically observed decision-making behavior more accurately than expected utility theory, and can explain a wide range of experimentally documented behavioral biases, including the status quo bias and endowment effects. One of the best-known achievements of behavioral economics, prospect theory led to the award of the 2002 Nobel prize in economics to Kahneman for his "... contributions to behavioral economics, in particular the development of cumulative prospect theory".

Yet the literature on mechanism design almost uniformly models agents as making choices according to the tenets of the classical expected utility theory. While the expected utility model may be an accurate description of choice-making for the applications addressed by classical mechanism design—such as auction theory applied to large firms which are possibly exempt from these behavioral biases, are the decision-making agents—a number of online systems that are the subject of several newer applications of mechanism design, such as crowdsourcing and labor markets, online auctions with small bidders, or peer-to-peer economies, depend on the decisions of individual agents to whom these behavioral biases do apply. What consequences might such biases in decision-making have for the design of incentives in these environments?

In this paper, we explore this idea of ‘behavioral’ design—how departures from standard economic models of agent behavior affect mechanism design—in the context of principal-agent problems in online labor and crowdsourcing markets. Broadly speaking, principal-agent problems—whereby a principal derives value from the outputs produced by agents who undertake his task (see Hart and Holmstrom [1987] for a survey)—requires analyzing how agents make choices under uncertainty: since an agent’s payoff depends, in general, on the (non-deterministic) quality of the output she produces as well as possibly that of other workers, the incentive effects of any particular mechanism or contract that the principal might use, and therefore the equilibrium value to the principal, depends on how agents make choices in the face of this uncertainty. We ask the question of whether a principal might want to choose incentive structures that are fundamentally different depending on whether agents behave according to prospect theory, or expected utility theory, in the setting of online labor and crowdsourcing markets, which we describe next.

1.1. Online labor and crowdsourcing markets

Jobs—of all kinds—are going online, with platforms for work ranging from microtasks such as Amazon Mechanical Turk (AMT or MTurk) to more substantial jobs including conventional ‘desk jobs’ such as software and data entry on online labor markets like oDesk, to research and development on platforms hosting crowdsourcing contests such as Kaggle and TopCoder. While there are, naturally, many similarities between online and traditional ‘offline’ employment, there are also some fundamental differences between the structure and organization of online and offline labor markets which point to interesting new economic questions.

A basic difference between online and offline labor markets comes from the short-term, ‘per-task’ nature of employment in online platforms, in contrast with longer-term employment in traditional labor settings where firms typically hire workers for more than just one task done once. The movement of labor online therefore means that a principal with a demand for workers can choose to use different platforms—and correspondingly different contract structures—for each task. The problem of which incentive scheme is most effective can then be addressed separately for individual tasks,

without being constrained by an organizational structure and existing employee assets of the firm which needs the task completed.¹

Workers' choice sets: Endogenous participation and exogenous quality. A second fundamental difference between offline and online environments comes from what workers strategize about. While the central issue in most of the principal-agent literature centers around effort elicitation—providing the right incentives for workers to elicit high effort (as a means to high output quality) in situations with imperfectly observable effort, incentives for *participation*, rather than quality, are the central issue in several online labor and crowdsourcing settings, for the following reasons:

- *Experimental evidence of participation-only choices:* A pattern discovered by multiple experimental studies investigating worker behavior in online crowdsourcing platforms is that a change in the level of offered incentives affects participation rates, but not the degree of effort that workers exert conditional on participation. As one example, Mason and Watts [2009] investigate the effect of financial compensation on performance in an online labor market for microtasks, and find in two experiments conducted on Amazon's Mechanical Turk (AMT) that increased financial incentives increase the quantity, but *not* the quality, of work performed by participants. In a very different kind of crowdsourcing task—answering questions in online Q&A forums—Jeon et al. [2010] investigate price as a predictor of quality in a field experiment on Google Answers, and also find the same pattern, namely that the prize offered to winning answers affects participation (or quantity), but not the quality of elicited answers. Both studies, therefore, display the same pattern of endogenous participation and exogenous quality, where increased pay increases how many users are willing to undertake the task, but not how much a user who has chosen to participate works at the task. Such behavior, possibly related to intrinsic motivation amongst workers for the task, arises in a range of crowdsourcing platforms ranging from microtasks to contests, and is summarized as a 'design claim' by Kraut and Resnick [2012] in their text on designing online communities: "With task-contingent rewards for small, discrete tasks, larger rewards will motivate people to take on tasks, but will not motivate higher effort on accepted tasks."
- *Tools for monitoring:* A different reason that effort (conditional on participation) may not be well-modeled as a strategic choice in online crowdwork platforms is the presence of monitoring tools that allow an employer to remotely observe and monitor workers. Virtual office applications for monitoring of workers, for example allowing employers to observe timed screen shots of workers' computers, are offered by various online labor platforms (including oDesk); these platforms also allow employers to make payments for work contingent on contractors performing the work while logged in to these virtual offices and monitored through regular screen shots and activity logs. With such monitoring software, the central issue of non-observability of effort driving the principal-agent problem disappears,² and the primary incentive question becomes one of incentivizing the appropriate number of workers to undertake the task.³

Consider, therefore, a principal with a single task that he wants to complete by hiring workers via an online platform. The principal has many options available in terms of what incentive structures he can use, such as those supported by various online crowdsourcing platforms—fixed-price contracts on platforms such as Amazon Mechanical Turk (MTurk) or their analog with price discovery, an auction to determine the payment for the task as on oDesk, as well as contests of various kinds on platforms like TopCoder or Kaggle. How should the principal choose between these various kinds of incentive structures—and does the answer to his question change depend on whether his population

¹In contrast, the decision of a firm with long-term employees between, say, whether to run an open-entry contest or simply assign some of its employees to work on a particular task, may depend on the skill set and size of its existing workforce.

²While such software can only monitor a worker's screen as opposed to her actual mental effort, the per-task nature of employment in online labor markets (such as oDesk) means that a worker who has undertaken a task might have strong incentives to put in her best effort to maximize her chances of future employment, with the same or a different requester.

³We do not model the issue of competition amongst principals which is beyond the scope of this paper; as we will shortly see, the principal-agent problem for incentivizing participation (as with effort) is already non-trivial with just one principal.

of potential workers behaves according to the classical model of expected utility theory, or whether they deviate, behaving instead according to prospect theory?

1.2. Our contributions

We investigate the question of whether behavioral biases in decision-making—as captured by prospect theory—can alter the answer to a principal’s incentive design problem in markets for on-line labor and crowdsourcing. We consider a simple model with a principal with a single task and a population of potential workers or agents, each of whom incurs an opportunity cost c to undertaking the task. Agents make endogenous participation choices but quality is exogenous (see §1.1), with each agent who strategically decides to undertake the task producing output with quality $q_i \sim F(q)$ at a cost c . To allow addressing the broader question of what kind of incentive contract the principal might prefer—as opposed to optimizing within a specific class of contracts (as, for instance, in an optimal contest design problem)—we use a general model for contracts encompassing a broad variety of incentive schemes, including those observed on major crowdsourcing platforms such as fixed prices to a fixed number of workers (MTurk), base payments with bonuses (oDesk), to open-entry tournaments with prizes to a small number of contestants (Kaggle, Topcoder).

We first consider a population of agents that make choices according to expected utility theory (§3). For such agents, we show in Theorem 3.1 that the *optimal* contract—for any increasing objective of the principal—always takes the form of an output-independent, fixed payment to each agent he employs: while the optimal *number* of agents to employ will depend on the specific parameters of the problem, the *structure* of payments does not. We then consider agents whose decision-making is described by prospect theory (§4), which models the empirically observed loss aversion via asymmetric value function $u(\cdot)$ over losses and gains and inaccurate weighting of probabilities via and, of particular significance to the incentive design question, a non-linear decision weight function $\pi(p)$ mapping probabilities to decision weights. To address the question of whether a principal would ever choose a fundamentally different kind of contract for agents with such prospect theory preferences, we analyze and compare the incentives created by contests—which are widely used in a variety of crowdsourcing applications and hosted by prominent crowdsourcing platforms such as Kaggle and TopCoder—against fixed payment contracts.

Contests are a very different kind of contract than fixed-payment schemes, in a qualitative sense: an agent’s reward in a contest is strongly output-contingent, depending not only on her own output quality q_i but also on the draws of all other participants in the contest. So when might a contest, which is inherently a riskier prospect, provide stronger incentives for participation than the safer fixed-payment contract with its output-independent payoffs? If agents used actual outcome probabilities as their decision weights, Theorem 3.1 can be extended to show that a principal would never use output-contingent payments for any risk-neutral or risk-averse population of agents, no matter what the nature of his task. With prospect theory preferences, however, a small chance of winning a large prize might contribute more than its ‘true’ share of utility due to the overweighting of small probabilities, potentially creating a larger perceived payoff for the same expected payout to the principal.

Despite the overweighting of small probabilities, however, a contest nonetheless does *not* always dominate a fixed-payment scheme: if the prize from winning the contest is not much larger than the cost of undertaking the associated task, few enough agents participate in equilibrium so that the probability of winning the prize becomes too large to be overweighted. This occurs even with risk-neutral agents (Theorem 4.2); when agents display risk and loss aversion, the overweighting of probabilities—favoring the contest—needs to also compensate for this aversion to uncertainty if a contest is to dominate a fixed-payment contract. We first derive a necessary and sufficient condition (Theorem 4.1) under which a contest will dominate the fixed-payment contract for general preferences (u, π) . Our first application of this condition, Theorem 4.3, shows when this condition *cannot* be satisfied: we show that, for any prospect theory preferences, a principal who does not have a large enough budget to spend on the task should not conduct a contest, even for agents who

overweight small probabilities—for such small prizes, risk and loss aversion beat out the benefits from overweighting of small probabilities.

More specific structure on the preference functions is required to say when the condition in Theorem 4.1 *can* be satisfied, *i.e.*, when a contest can do better than fixed-payment contracts; for this, we use the functional forms for u and π from the literature on econometric estimation of the cumulative prospect theory model [Tversky and Kahneman 1992; Bruhin et al. 2010]. Since these functional forms are based on fitting the prospect theory model to extensive experimental data, they are the most natural candidates for our need for specific preference functions to evaluate (3), yielding the best answer we can hope to have to our question—without conducting a measurement of the functions u and π in the specific marketplace of interest—for ‘real’ population preferences u, π . Our main results in §4.2, Theorem 4.4 and Theorem 4.5, together with the estimated parameter values from Tversky and Kahneman [1992]; Bruhin et al. [2010], provide an affirmative answer to our central question: to the extent that these parameters indeed describe the decision-making behavior of agents in online crowdsourcing environments, a principal who values the output from crowdsourcing his task sufficiently highly (compared to the cost to a worker to produce that output) might indeed choose a different kind of mechanism, and do better as a result, if he accounts for deviations from the standard economic model of expected utilities that are typically used in theoretical design.

1.3. Related work

A vast literature, too large to properly cite here, studies the optimal design of labor contracts, as well as rank-order contests as incentive schemes for the procurement of labor, beginning with Lazear and Rosen [1981]; Green and Stokey [1983]. The subset of this literature that is most relevant to our work concerns comparing amongst mechanisms for the procurement of innovation [Che and Gale 2003; Fullerton et al. 2002; Schottner 2008]: this work, like ours, compares different kinds of incentive structures as opposed to deriving the optimal mechanism of a *given* kind (such as the optimal contest or the optimal bonus-payment scheme). The most relevant of these is the beautiful work of Che and Gale [2003], which derives the optimal mechanism for procuring innovations in a very general class of market structures. However, while asking the same broad question about what kind of incentive structure (rather than specifics of which mechanism of a certain kind) is optimal, our work and Che and Gale [2003] differ in two fundamental ways. First, the agents in Che and Gale [2003] are standard expected utility maximizers, while we allow for agents to have more general prospect theory preferences, in keeping with our primary interest in contrasting the two models of decision-making. Second, agents in Che and Gale [2003] make strategic effort choices, unlike in our setting where quality is exogenous and only participation is a strategic choice (§1.1). These differences in choice sets and decision-making behavior, in addition to consequential differences in information structure between the motivating environments, fundamentally alter the nature of the incentive problem across the two settings.

There is also an extensive body of work on contest design, studying several variants of the problem of optimally designing a contest’s reward structure—the number and amounts of the prizes for each rank is the most common version—in a variety of models including, for example, heterogeneous and homogeneous agent populations, risk-neutral and risk-averse preferences, and non-monetary rewards, as well as various objectives and constraints for the principal running the contest. The work closest to ours from this literature is Ghosh and Kleinberg [2014], which considers contest design in a similar model with strategic participation decisions but exogenous quality (although a somewhat different information structure), and asks how this departure from the strategic effort choice model that is typical in the contest design literature can change the structure of the optimal contest design problem. The primary difference between our work and the contest design literature, in addition to the differences in expected utility versus prospect theory preferences and whether quality is strategically chosen or non-strategic,⁴ is the question addressed: the contest design liter-

⁴With the exception noted above, agents are almost uniformly modeled as making strategic effort choices that affect their outputs quality.

ature takes the incentive structure of a rank-based contest as given, assuming that the principal (for a variety of well-justified reasons) will use a contest mechanism, and asks what is the best contest mechanism he can use. In contrast, we do not assume that the principal is bound to a contest mechanism at all, but rather can choose amongst a much more general space of contracts that include, but are not limited to, rank-order contests. Thus, while we ask about choosing amongst a much richer class of mechanisms, in a simpler model, the contest design literature asks about choosing amongst a more restricted set of mechanisms, but in a richer model (albeit without prospect theory preferences).

Finally, while prospect theory has been extensively studied in the behavioral economics literature, it has seen relatively little application in theoretical work. The settings and models addressed in these papers, as well as the nature of the questions posed, in the few exceptions [de Meza and Webb 2007; Kanbur et al. 2008; Nakajima 2011; Chang et al. 2014] are very different from our work.

2. MODEL

We begin by describing a simple model to address the question of how a principal using an online labor or crowdwork platform to recruit workers should choose between various kinds of incentive structures—such as fixed prices or base payments with bonuses as on MTurk and oDesk, or crowdsourcing contests as on Kaggle or Topcoder—and how this answer depends on workers decision-making behavior. To allow addressing the broader question of what kind of incentive contract the principal might prefer—as opposed to optimizing within a specific class of contracts (as, for instance, in an optimal contest design problem)—we use a very general model for contracts encompassing a broad variety of incentive schemes, including those supported by the largest online crowdsourcing platforms. The principal’s choice of incentive scheme affects both his total expected payout and the extent of participation by agents, which are determined in an equilibrium of the participation game induced by the contract. The model, described in detail below, allows us to ask what kind of mechanism the principal should choose, both for classical expected utility maximizing workers and for workers who deviate from the expected utility model, behaving instead according to prospect theory preferences.

2.1. Agents

There is a large pool of agents, or workers, indexed $i = 1, 2, \dots, N$, each of whom strategically decides whether or not she would like to work on the single task posed by a principal. As discussed in §1.1, we consider scenarios with exogenous quality and endogenous participation, *i.e.*, where agents make strategic participation choices, but do not strategize over the quality of their output. We model this by assuming that qualities are random draws from a fixed distribution: an agent who undertakes the task produces output of quality $q_i \in \mathbb{R}_{++}^1$ where the qualities q_i are IID draws from a distribution⁵ with CDF $F(q)$. The realization of the quality draw q_i becomes known to the agent (and the principal) only *after* the agent incurs the cost associated with undertaking the task, corresponding to typical crowdwork scenarios where uncertainty in an agent’s output quality is resolved only after the agent works on the task (for example, in microtasks such as image labeling, as well as with larger tasks such as designing algorithms in data mining contests).⁶ We consider

⁵The assumption that all workers’ qualities are drawn from the same distribution F says that the principal or requester does not have a priori information predicting which workers will produce better outputs on the task. If such information is available, the assumption still remains relevant if we suppose that the requester chooses amongst a subset of workers all of whom have the same ‘best’ distribution. Note also that qualities are non-zero for participating workers; we reserve 0 for workers who do not participate.

⁶While appropriate for our setting, we note that this assumption differs from that in some other prior work, for example, where an agent’s quality is either the result of her strategic effort choice (as in Che and Gale [2003], where c and q are essentially ‘revealed simultaneously when the agent chooses her effort), or in situations where the cost c models the cost of submitting existing output of known quality to the principal (such as an entry fee in a contest) as in Ghosh and Kleinberg [2014], resulting in corresponding differences in the incentive structure.

agents who are ex-ante homogeneous, so that each agent has an opportunity cost $c_i = c$ of working on the task (we address heterogeneous costs in the appendix of the full version of the paper).

Homogeneity amongst workers, while not without loss of generality, is nonetheless a reasonable assumption in a variety of crowdsourcing environments. For instance, consider ‘routine’ tasks such as image labeling on MTurk or data entry tasks on online labor platforms. If c models the opportunity cost of time spent on a task (e.g. the price offered by other requesters for similar tasks in the market), and routine tasks take all workers essentially the same time, workers would incur homogeneous costs c . Similarly, if c models the effort cost for a task, and the routineness of the task is such that workers do not have vastly differing abilities or skills for the task, it is again reasonable to suppose that c is the same (to a first order) for all workers. For non-routine tasks (such as innovation or development) also, homogeneity amongst the population of potential participants can apply, to a first approximation, for at least one of two reasons: *self-selection*, such as in contest environments (for example, programmers with similar levels of expertise or graphic designers with similar skill levels), as well as *filtering* or selection by the requester (the principal), resulting in workers with similar ability or expertise levels, and therefore similar costs, either opportunity or effort costs, to undertaking the task (for example, AMT supports a feature whereby a requester can specify ‘qualifications’—including customizable, self-defined ones—which a worker must satisfy in order to be considered for the task⁷). In addition to the appendix, see also §5 for a more general discussion on heterogeneity in worker populations.

2.2. Preferences and decision making: Expected utility and prospect theory

A worker in a principal-agent problem (as in the crowdsourcing scenarios we study) faces uncertain payoffs, since her reward from undertaking the task may depend on the non-deterministic quality of her output as well as those of other workers employed by the principal. The worker therefore faces a problem of decision making under uncertainty, the classical economic model for which is expected utility theory. Suppose a worker who undertakes the task will receive payoffs $x_k, k = 1, \dots, K$, with corresponding probabilities p_k ; such a $(p_1, x_1, \dots, p_K, x_K)$ tuple is called a prospect or gamble.⁸ According to the expected utility model, an agent ascribes a utility $u(x_k)$ to each possible payoff x_k (where u is typically assumed to be concave, to model risk aversion), and makes decisions by evaluating the expected utility $\sum_k p_k u(x_k)$ from the prospect. We assume that the set of possible payoffs contains the payoff $x_k = 0$, and that non-participation in the task yields this zero payoff.

In their seminal work on decision making under risk, Kahneman and Tversky [1979] presented experimental evidence of choice-making behavior that is irreconcilable with this classical expected utility model, and proposed an alternative model—prospect theory—that can explain systematic deviations from expected utility theory such as inaccuracy in weighting payoffs by their probabilities, and greater disutility for losses compared to utilities from a gain of the same amount. In this paper, we will be interested in the question of whether these deviations from the standard model can have an effect on the outcome of the mechanism design problem faced by the principal.

We use the following abstract model, which includes expected utility theory as a special case, to describe agents with prospect theory preferences in our setting. Given a task that yields uncertain payoffs x_k with corresponding probabilities p_k , an agent decides whether to undertake the task by evaluating its net benefit according to prospect theory preferences [Kahneman and Tversky 1979] as

$$\sum_k \pi(p_k) u(x_k), \quad (1)$$

⁷Quoting from AMT: “A Qualification is a property of a Worker that represents a Worker’s skill, ability or reputation. You can use Qualifications to control which Workers can perform your HITs. A HIT can have Qualification requirements that a Worker’s Qualifications must meet before the Worker is allowed to accept the HIT. . . . You can create and maintain your own Qualifications using the web service API.”

⁸Note that these payoffs x_k and their probabilities p_k might be determined in an equilibrium of the game corresponding to a particular contracting mechanism used by the principal; here we are interested in describing how an agent evaluates an uncertain proposition given x_k and p_k , rather than where these values come from.

where $u(\cdot)$ is the value function over the payoffs x_k , and the $\pi(\cdot)$ are decision weights that map the probabilities p_k of each outcome into the weights with which they contribute to the perceived total utility from the gamble. Without loss of generality, we normalize the utility function so that the utility of current wealth, a payoff of 0 is 0.⁹

The standard model of expected utility maximizing agents is obtained by setting the decision weight function π to be the identity function ($\pi(p) = p$ for all $p \in [0, 1]$), and the utility function u to be strictly concave to model risk aversion. *Prospect theory* [Kahneman and Tversky 1979; Tversky and Kahneman 1992] deviates in both the decision function π and the structure of u . First, prospect theory says that agents do not correctly weigh uncertain outcomes, except at the extremes of absolute impossibility ($p = 0$) and absolute certainty ($p = 1$): rather than weight payoffs via the identity function $\pi(p) = p$, they tend to act as if they overestimate the likelihood of very low probability events or equivalently overweight unlikely outcomes ($\pi(p) > p$ for small p), and underweight highly likely events that are not completely certain (corresponding to $\pi(p) < p$ for large p). We abstract this property of the decision weight function π from prospect theory as follows:

There exists a $p^* \in (0, 1)$ such that $\pi(p^*) = p^*$, $\pi(p) \geq p$ for $p < p^*$, and $\pi(p) \leq p$ for $p > p^*$.

Second, prospect theory models *loss aversion*—the observation that a loss of x dollars typically causes more pain than the pleasure from gaining the same amount x —via asymmetry around 0 in the value function u over payoffs.¹⁰ In prospect theory, the function u , which is concave for positive payoffs, convex for negative payoffs, and satisfies $u(0) = 0$, has the property $-u(-x) > u(x)$ for all $x > 0$.

In our context, a worker with prospect theory preferences will decide whether or not to participate in a task with opportunity cost c and offering payoffs x_k with probabilities p_k for $k = 1, \dots, K$ (and zero payoff with the remaining probability $p_0 = 1 - \sum p_k$) by comparing her perceived total utility from this gamble against the payoff from not-participating. Deciding to participate results in a cost of c with probability 1, contributing $\pi(1)u(-c)$ to the perceived utility, and a benefit of x_k with probability p_k , each contributing $\pi(p_k)u(x_k)$, while the payoff from not participating is 0. Using $u(0) = 0$, and $\pi(1) = 1$, an agent will therefore decide to undertake the task if and only if:

$$\sum_k \pi(p_k)u(x_k) + u(-c) \geq 0. \quad (2)$$

Of course, the payoffs x_k and their probabilities p_k are typically not exogenously specified, but rather arise as an equilibrium of the game induced by the principal’s chosen contracting mechanism; we discuss this next.

2.3. Contracts

We consider a principal, or firm—a ‘requester’—who has a task or project that he wishes to crowdsource amongst the pool of workers on a crowdsourcing platform (we use the term firm and principal interchangeably throughout the paper). The principal can choose how many workers, amongst the set of workers who are willing to undertake his task, he would like to employ, as well as the (possibly output-dependent) payments that he will make to each of these workers. We are interested in the question of what form of contract a principal should use, motivated by the observation that many different kinds of contracts, ranging from fixed wage payments to a fixed number of workers, to open-entry tournaments with prizes to a very small number of contestants, are observed across various crowdsourcing platforms. To allow addressing this question (rather

⁹This normalization is free in the sense that the normalized and un-normalized utility function represent the same preferences over gambles, but the fact that current wealth has utility 0 has no cardinal meaning. In prospect theory, however, “utility” is defined over gains and losses and so making the starting point 0 does have meaning.

¹⁰The value function u in prospect theory is defined with respect to the reference point of the agent’s current wealth *i.e.* it takes as argument only gains or losses (the possible payoffs x_k and loss c) relative to the reference point

than the more restricted question asking what is the optimal mechanism of a certain form, for instance the optimal contest), we will first need an adequately general definition of mechanisms.

Mechanisms. Consider a task with a set of N potential workers, each of whom decides whether or not to undertake the task. Workers who undertake the task produce outputs with qualities $q_i \sim F(q)$; if a worker does not participate, we denote her output quality by 0. Let $[\mathbf{q}]$ denote the vector of all N qualities. A *mechanism* \mathcal{M} is specified by the schedule of payments $(w_1([\mathbf{q}]), w_2([\mathbf{q}]), \dots, w_N([\mathbf{q}]))$ that the principal will make to each of the $i = 1, \dots, N$ workers, for all possible values of the vector of output qualities $[\mathbf{q}]$. (If workers are ex-ante homogeneous (as we assume throughout, except in the appendix), the payment functions w_i need not depend on the ‘name’ of a worker (although they may still depend on the worker’s output quality q_i .)

Note that this is a general abstraction of a contract, encompassing a number of observed market mechanisms:

- (1) *Fixed payment contracts to a fixed number of workers.* A principal may specify that he requires N_{FP} workers, and that he will pay a uniform price w to every worker (up to N_{FP}) that completes his task. As an example, requesters on Amazon Mechanical Turk can specify the number of HITs (Human Intelligence Tasks) available for a given job, as well as the price for the HIT that will be paid to each worker upon its completion. This corresponds to $w_i([\mathbf{q}]) = w$ if i is amongst the set of (up to) N_{FP} workers employed by the principal, and $w_i([\mathbf{q}]) = 0$ otherwise.
- (2) *Fixed payments with bonuses.* The payment specified for each participating worker may be made contingent upon the quality of her output. For example, a contractor on oDesk may specify that he will hire one worker who receives some pre-determined base payment for the task plus possibly an additional, quality-dependent bonus. Such fixed-payment-plus-bonus schemes are captured by setting $w_i([\mathbf{q}]) = w_i^0 + w_i(q_i)$ for every agent i employed by the principal, and $w_i([\mathbf{q}]) = 0$ for all other agents.
- (3) *Rank-order contests.* Both previous schemes make payments to workers that do not depend on the quality of the output produced by other agents working on the task: the first scheme is completely agnostic to the realized qualities, while the second scheme bases an agent’s payment only on the quality of her own output. A commonly observed mechanism in crowdsourcing scenarios is a rank-order contest, where the principal announces a schedule of prizes to be awarded to the agents with the k highest-ranked outputs. For instance, Topcoder runs competitions where the competitor who submits the best code for an algorithm design problem receives a pre-specified prize while all other competitors win nothing; Kaggle hosts contests with a wide range of prize structures—winner-take-all contests, equal or unequal prizes to the top three submissions, and so on. Contests are captured by setting $w_i([\mathbf{q}]) = w^{\pi(i)}$, where $\pi(i)$ is the rank of q_i in $[\mathbf{q}]$, and $w^j \geq 0$ is the pre-announced prize for the j -th ranked entry.
- (4) *Cardinal contests.* The payment to an agent may depend both on her rank $\pi(i)$ and her absolute output q_i —such cardinal contests [Ghosh and Hummel 2015] are growing in popularity due to contest environments with measurable output qualities. For instance, some data mining contests on Kaggle specify the prizes that will be awarded to the top k -contestants provided their algorithm’s performance exceeds a certain baseline or threshold \bar{q} ; this corresponds to a contract with $w_i([\mathbf{q}]) = w^{\pi(i)} 1_{q_i \geq \bar{q}}$.

Equilibrium participation $N_{\mathcal{M}}$. Each potential worker in the market needs to make a decision about whether she would like to undertake the task or not, given the mechanism \mathcal{M} that will be used by the principal to make payments. At the time she makes her decision, a worker knows the cost c associated with the task, the preference parameters $\pi(\cdot)$ and $u(\cdot)$ and the distribution F from which her (and other workers’) qualities will be drawn. However, she does not know the actual realization of her output quality q_i (or that of other agents)—this is revealed (both to her and the principal) only after incurring the cost c to perform the task, as in typical crowdsourcing scenarios where the quality uncertainty about output is resolved only after a worker attempts the task. (Note that this means that

while the payment may depend on the realized qualities q_i , any pre-selection by the principal of which workers to hire—*i.e.*, choosing which subset of agents will undertake the task (incurring the cost of c), for example as in (1) above—cannot use information about the realized qualities q_i , since this information is unavailable to the principal or agents until the agent undertakes the task.)

A choice of mechanism \mathcal{M} by the principal induces a game; recall that workers only make strategic participation choices in this game. Given a mechanism \mathcal{M} the number¹¹ of workers $N_{\mathcal{M}}$ that decide to undertake a task is determined *endogenously* as a Nash equilibrium in the participation game induced by \mathcal{M} as follows. Workers’ participation decisions (and the distribution F) induce a distribution on the vector of output qualities $[\mathbf{q}] = (q_1, \dots, q_N)$ (with non-participation by worker i modeled as $q_i = 0$), which in turn induces a distribution on the payments $w[\mathbf{q}]$ to each worker. Given this distribution on payoffs induced by all workers’ participation decisions (and the qualities F), an equilibrium is one in which all participants’ value from this prospect (§2.2) is greater than or equal to 0, and would be less than or equal to 0 to non-participants if any one of them chose to participate. Note that as each worker makes only a yes-no participation decision, there is at least one equilibrium. To simplify the discussion we will ignore the discreteness of the number of workers and simply assume that in any equilibrium the payoff to the marginal worker is 0.

2.4. Principal

Different kinds of tasks benefit from multiple workers in different ways: for instance, a requester whose task has little or no uncertainty about quality may place little value on having more than one worker to complete this routine task. Alternatively, a requester whose task has highly variable qualities, and who values high quality, may want to have many workers attempt the task. We use the function V to model how the principal benefits from the output of multiple workers: the value to the principal from an ordered vector¹² of qualities $(q[1], \dots, q[N])$ is $V(q[1], \dots, q[N])$. We assume that $V(0, \dots, 0) = 0$, *i.e.*, the principal derives no value if no worker undertakes his task. Since workers’ qualities all come from the same distribution F , the expected value to the principal depends only on the *number* of workers incentivized to participate in \mathcal{M} , *i.e.*, $EU_p(\mathcal{M}) = V(N_{\mathcal{M}})$, and of course the distribution F . In addition to the mechanism itself which plays an important role in determining the equilibrium participation $N_{\mathcal{M}}$, note that $N_{\mathcal{M}}$ will, in general, depend on the exogenous parameters describing the population of workers: their costs c , the functions describing their preferences in decision making—the value function $u(\cdot)$ over payoffs, and the decision weights $\pi(p)$, as well as the distribution F from which output qualities are drawn.

Information structure. We assume, as is standard in the literature, that the principal knows the distribution F from which workers’ output qualities for his task are drawn and the parameters describing the decision-making preferences—the valuation functions u and decision weights π —of the worker population. We also assume that the cost c of undertaking the task is known to the principal: this is a reasonable assumption, for example, in markets where there are many tasks of the same kind (e.g. image labeling, or audio transcription tasks of a certain length, on MTurk), so that a requester can likely easily estimate the opportunity cost to workers for his task. In markets where the principal cannot accurately gauge the cost for her task (for instance, if every task in the market is different), the value of c can be elicited by a price-discovery process such as an auction (indeed, one observes fixed prices without the need for auctions for price discovery on MTurk while oDesk supports, and requesters do use, auctions where workers can bid their price for a job). In either case, the price-discovery via the auction does not affect our arguments, which center around whether payments are contingent on output quality or not. This holds even if agents have heterogeneous costs

¹¹Note that with homogeneous workers, the relevant quantity is the equilibrium distribution of the number of participants.

¹²Note that this form of V implicitly assumes that the principal’s value depends only on the set of qualities produced, not on which worker produced which quality.

c_i , as long as there are not unknown correlations between the cost c_i to the worker and the quality q_i of the agent's output.¹³

Objective. For any mechanism \mathcal{M} , the principal can compute the expected payoff in equilibrium in the participation game induced by the contract \mathcal{M} . We consider a principal who has a fixed amount W to spend (for instance, a manager who has been given a total feasible spend to have the task completed, without concern for leftover money) and wants to maximize its value over all mechanisms \mathcal{M} that spend (no more than) W . An optimal contract is thus a contract \mathcal{M} costing no more than W that induces an $N_{\mathcal{M}}$ that maximizes $V(N_{\mathcal{M}})$ over all feasible contracts.

3. OPTIMAL CONTRACTS: EXPECTED UTILITY MAXIMIZERS

We begin by investigating the nature of the optimal contract that a principal should use to employ agents when the population of workers behaves according to the standard model of expected utility theory—namely, agents with decision weights that equal probabilities, $\pi(p) = p$, and risk-neutral or risk-averse preferences corresponding to concave functions $u(x)$ (risk neutrality corresponds to the special case of linear u). In such an environment—modeling a wide range of crowdsourcing markets—where agents make strategic participation choices about whether or not to participate in a task (but do not strategize over quality), what kind of contract \mathcal{M} —for instance, fixed prices (as in MTurk), fixed payments with bonuses (as in some tasks on MTurk or oDesk), or contests as in several crowdsourcing platforms (Kaggle, TopCoder)—is optimal for the principal? The answer to this question could potentially depend on a number of parameters, such as the dependence of the principal's utility on the number of workers $V(n)$, the scale of the principal's utility V compared to workers' costs c , the spread in workers' output qualities F , and so on.

Our main result in this section, Theorem 3.1, shows that the optimal contract always takes the form of a fixed price mechanism, where the principal pays an optimally chosen number of agents a fixed price¹⁴ independent of output quality. The proof, presented in the appendix, hinges on agents' not being risk-seeking, and proceeds by arguing that for *any* total expected payout W the principal might make, that payout W incentivizes the highest participation when it is disbursed as fixed payments rather than via any output-contingent scheme, if the agent population has risk-averse or risk-neutral preferences.

THEOREM 3.1. *Consider a market where agents are expected utility maximizers with concave valuation functions, $u(\cdot)$, over payoffs. Suppose agents have an opportunity cost c to participation and qualities q_i exogenously drawn from a distribution F . Then, for any increasing value function $V(n) = V(q_1, \dots, q_n)$, any c, F , and any equilibrium of any mechanism \mathcal{M} there is an equilibrium of a fixed-price mechanism in which:*

- Each worker makes the same participation decision as in the given equilibrium of \mathcal{M} ,
- The firm's value is at least as large as in the given equilibrium of \mathcal{M} .

This result is striking when viewed in light of the wide range of contract structures that our model allows the principal to choose from—for instance, making payments output-dependent to hedge against the possibility of paying too much for poor output could arguably improve utility over a fixed-price scheme; as another reasonable hypothesis, using contests to get away with paying a few agents while benefiting from the efforts of a larger number of agents. Nonetheless, the theorem shows that the principal is always best off making payments that are not output-contingent, no matter what his value function V or the agent parameters c, F, u are; those parameters only affect the *number* of agents that the principal would optimally hire. (See the appendix for the analog of this result with heterogeneous agents, which retains the same central feature of pre-specified payments that are not contingent on an agent's absolute or relative output quality.)

¹³The case of such correlations leads to a significantly more complex two-dimensional mechanism design problem, which is both beyond the scope of the paper and not central to our primary interest about the differences resulting from prospect theory versus expected utility behavior; see §5.

¹⁴This price is c ; if the value of c is private information to agents it can be elicited using an auction mechanism

We note that there is a conceptual similarity between this argument and a basic result in the principal-agent literature (albeit in a different model) regarding the optimal mechanism for a principal facing a single agent if the principal could directly observe and reward effort, setting the comparison point for a principal facing an agent whose strategic, unobservable, effort choice affects the output she produces. Theorem 3.1 above also sets a comparison point for us, although in quite a different sense: it tells us that the principal is always best off making payments that are not output-contingent, no matter what his value function V or the agent parameters c, F, u are (which only affect the number of agents to optimally hire), as long as he faces expected utility maximizing agents. This sets the stage for asking whether prospect theory preferences might lead a principal to prefer to use commonly observed alternative contract structures—contests in particular—in crowdsourcing markets which are reasonably described by our model.

4. CONTRACT DESIGN: PROSPECT THEORY PREFERENCES

We now ask whether behavioral biases can indeed matter to the incentive design problem faced by a principal: would a principal ever choose a fundamentally different contract for agents with prospect theory preferences than that for expected utility maximizing agents? In this section, we address this question by comparing the incentives created by contests, which are widely used in a variety of crowdsourcing applications, against fixed payment contracts which were shown in §3 to be optimal for expected utility maximizing agents. Specifically, the question we ask is the following. Consider a given population of agents with preferences described by some value and decision weight functions (u, π) , and a principal with a task with cost c and distribution F of workers' output qualities. In this (c, F, u, π) environment, are there any values of W —the total amount the principal wants to spend on his task—for which he is better off using a winner-take-all contest with prize W , where the agent who produces the highest-quality output amongst all participants receives W and all other contestants receive nothing, instead of a fixed-payment contract?

We address this question, instead of the optimal contract design problem as in §3, for two reasons.¹⁵ First, it is reasonable to restrict attention to contract structures that a principal crowdsourcing his task can realistically offer to workers, which are those supported by major online crowdsourcing platforms with large worker bases (online labor markets like MTurk and oDesk, or contest platforms such as Kaggle and TopCoder). While all of these platforms offer significant flexibility to requesters about the specific terms of a contract (for instance, the actual prices for labor or the number of workers employable per task on oDesk or MTurk, or the prize amount and number of prizes offered on, say, Kaggle), the structure of supported contracts most commonly take the form of either fixed payments (possibly with performance-based bonuses), or contests (we briefly discuss contests with more general prize structures than the winner-take-all contest in the appendix). Second and more important, we choose the simplest set of results that illustrate the central point we want to make, namely that behavioral biases—deviations from standard economic models of decision-making—can indeed make a fundamental difference to the principal's choice of mechanism in some realistic environment, which purpose is served by the outcome of this comparison.

4.1. Contests versus fixed payments: Preliminaries

Contests are a very different kind of contract than fixed-payment schemes, in a qualitative sense: in contrast with the fixed-payment scheme from §3, an agent's reward in a contest is output-contingent, depending not only on her own output quality q_i but also on the draws of all other participants in the contest. So when might a contest, which is inherently a riskier prospect, provide stronger incentives for participation than the fixed-payment contract with its output-independent payoffs? If agents used actual outcome probabilities as their decision weights, Theorem 3.1 shows that a principal would never use output-contingent payments for any risk-neutral or risk-averse population of agents, no

¹⁵In addition to these reasons, note that the structure of the optimal contract can possibly be quite sensitive to specific parameters describing both the population of workers (c, u, π) and the nature of the principal's task (F, V) when agents are not expected utility maximizers.

matter what the nature of his task.¹⁶ With prospect theory preferences, however, a small chance of winning a large prize might contribute more than its ‘true’ share of utility due to the overweighting of small probabilities, favoring a contest.

To understand the intuition behind why contests might provide stronger incentives with overweighting of small probabilities, assume for a moment that an agent will ‘win’ the prize W if her quality draw beats an exogenous threshold q^* (imagine, for example, that the principal derives a value W if the output quality is greater equal the current state of the art q^* , and no value otherwise). This event $q_i \geq q^*$ has probability $\epsilon = 1 - F(q^*)$, and the perceived payoff to the agent from the contest is $u(W)\pi(\epsilon)$. A different payment scheme that has the same expected payout to the principal is the following: rather than wait to see if the output quality exceeds q^* before making the payment, the (risk-neutral) principal commits to paying his expected value, which is $W\epsilon$, yielding payoff $u(W\epsilon)$ to the worker. To develop intuition, imagine that the agent is risk-neutral so that $u(x) = \alpha x$; the agent then perceives respective payoffs $\alpha W\pi(\epsilon)$ and $\alpha W\epsilon$ from the output-contingent contest and fixed-payment schemes respectively. Now if the chance of winning the contest $\epsilon = 1 - F(q^*)$ is small enough for the overweighting of small probabilities to become relevant, the agent overweights the outcome of winning W ($\pi(\epsilon) > \epsilon$), leading to a larger perceived benefit ($\alpha W\pi(\epsilon)$) from the contest than from the fixed-payment scheme ($\alpha W\epsilon$), and correspondingly stronger incentives for participation. Of course, this reasoning is oversimplified—it ignores risk aversion, corresponding to non-linear u ; more importantly, it ignores the fact that the threshold q^* that an agent’s output q_i needs to beat in a contest is not exogenous but rather is endogenously determined in equilibrium by the choices of all other workers: $q^*(N^*) = \max_{j=2,\dots,N^*} \{q_j\}$ where N^* is the equilibrium number of participants in the contest. However, it illustrates why overweighting of small probabilities can skew incentives in favor of ‘riskier’ gambles than ‘safer’ prospects, under the right circumstances.¹⁷

Even when agents use decision weights that overweight small probabilities, though, it will not always be the case that contests yield a more attractive incentive scheme than fixed-prices, because of risk aversion: the increased weighting of small probabilities by π —favoring contests where there is a small chance of winning a large prize—trades off against the risk and loss aversion in the value function u , which favors the smaller but certain payoffs awarded by a fixed-price scheme. This tradeoff between u and π in the prospect theory preferences will be central to whether the riskier, output-contingent, contest can ever yield stronger incentives than ‘safe’ fixed-payment contracts.

We first prove a theorem that tells us whether a contest will provide better incentives than a fixed payment scheme in terms of the value and decision weight functions u, π describing agents’ preferences. This result forms the basis for all our remaining results in this section.

THEOREM 4.1. *Consider a population of agents with value function $u(x)$ and decision weights $\pi(p)$, and a task with cost c . A winner-take-all contest with prize W dominates a fixed-payment contract with total payout W in this environment with preferences (u, π) if and only if W is such that*

$$\pi\left(\frac{c}{W}\right) > \frac{-u(-c)}{u(W)}. \quad (3)$$

Note that the proof of this result does not rely on any properties of the u and π functions except that the decision weights $\pi(p)$ are increasing in p ; it is merely a comparison between the numbers of participants for any preferences described by u and π .

We now want to understand ‘when’, *i.e.*, for what kinds of total payouts W and preferences u, π , this inequality can hold, since this can help us answer our question about whether behavioral biases can make a difference to incentive design: if there are (realistic) (u, π) such that this inequality

¹⁶If agents’ preferences display loss aversion $-u(-x) > u(x)$ in addition to risk aversion this only makes output-contingent payments less valuable to them reinforcing the attractiveness of a fixed payment contract.

¹⁷This same reasoning can be used to understand why an output-dependent bonus payment, while strictly dominated for risk-averse agents with linear weighting of probabilities, may similarly improve incentives if agents overweight small probabilities for the same expected payout to the principal.

does hold for some range of payouts W , the designer can indeed do better with a contest for such (u, π, W) than with the fixed-payment contract which is optimal if agents behave according to expected utility theory (Theorem 3.1).

Before discussing when (3) can hold, however, we point out first when it *cannot* hold. Suppose the decision weights π are linear ($\pi(p) = p$), so that agents weigh outcomes by their actual probabilities when computing their expected payoff. In this case, $\pi(\frac{c}{W}) = \frac{c}{W}$ but $\frac{c}{W} \leq \frac{-u(-c)}{u(W)}$ for preferences u with risk and loss aversion (as shown in the proof of Theorem 4.3), so that (3) cannot hold. Therefore, if agents accurately weigh outcomes by their probabilities, a designer will never—for any value of the total payout W —prefer to use a contest over a fixed-payment contract.¹⁸ The difference between probabilities and the corresponding decision weights—between p and $\pi(p)$ —in the prospect theory preferences must therefore be central to any possible difference in the principal’s choice of contract.

Risk-neutral agents. Recall that the overweighting of small probabilities and under-weighting of larger ones in prospect theory is modeled by saying that exists a $p^* \in (0, 1)$ such that $\pi(p) \geq p$ for $p < p^*$, $\pi(p^*) = p^*$, and $\pi(p) \leq p$ for $p > p^*$ (see §2). Our first result below shows that a contest does not always dominate a fixed-payment scheme despite the overweighting of small probabilities—which makes the high-risk contest seem a more favorable incentive scheme—even with risk-neutral agents: the reason is that if the total prize from the contest is not much larger than the cost of undertaking the contest task, few enough agents participate in equilibrium so that the probability of winning the prize becomes too large to be overweighted. Therefore, it is only when W/c is large enough that a contest can dominate a fixed-payment contract. The proof follows immediately from Theorem 4.1 by noting that if u is linear, $-u(-c)/u(W)$ is c/W .

THEOREM 4.2. *Suppose agents are risk-neutral, that is, u is linear. Then, a winner-take-all contest dominates fixed payment contracts if and only if the total payment $W > c/p^*$.*

Risk-averse and loss-averse agents. In general, of course, we would not expect agents to be risk-neutral. Prospect theory models u as concave on $x > 0$ to capture risk aversion, and $u(x) < -u(-x)$ to model loss aversion. Now the increased weighting of small probabilities by π has to additionally also compensate for the risk and loss aversion in the value function u , which favors the smaller—but certain—payoffs with fixed payments, if the contest is to dominate the fixed-payment contract. Unlike the case of linear value functions u corresponding to risk neutrality, the range of values W for which a contest dominates a fixed-payment contract need not, for an arbitrary concave preference u , be as nicely structured as an interval, and in fact, need not even be non-empty.

While obtaining a result about when a contest does dominate a fixed-payment contract (*i.e.*, when inequality (3) will be satisfied) requires more specific structure on the function u than concavity, we can give conditions under which a contest will not dominate a fixed-payment contract for any preferences u satisfying risk and loss aversion. Theorem 4.3 below says that when the total payment W is not large—specifically, not larger than c/p^* —contests cannot dominate fixed payment contracts. So a principal who does not have a large enough budget to spend on the task should not conduct a contest, even for agents with prospect theory preferences: for such small W , risk and loss aversion (concave u) beat out the benefits from overweighting of small probabilities ($\pi(p) > p$ for small p).

THEOREM 4.3. *For any prospect theory preferences (u, π) , a winner-take-all contest cannot dominate a fixed-payment contract for any total payout $W \leq c/p^*$.*

¹⁸Note that this is closely related to our claim from §3 that a fixed-payment contract (weakly) dominates all other forms of contracts for expected utility maximizing agents with risk-neutral or risk-averse valuations over payoffs, and is in fact a special case of that claim if there is no loss aversion in u .

4.2. Contests versus fixed payments: Decision weights from cumulative prospect theory

Recall that our primary question was whether a principal might choose a contract different from that which is optimal for expected-utility agents when facing a population of agents who demonstrate behavioral biases; to this end, Theorem 4.1 gives us a necessary and sufficient condition (3) under which a contest will dominate the fixed-payment contract. So far, we have seen when this condition *cannot* be satisfied for general prospect theory preferences u, π ; however, saying when (3) *is* satisfied requires more specific structure on u and π .

In this section, we use specific forms of the utility functions u and decision weights π from the literature on econometric estimation of the cumulative prospect theory model [Tversky and Kahneman 1992], a theoretical refinement of the prospect theory model in Kahneman and Tversky [1979]. Since the functional forms in this literature are based on fitting prospect theory preferences to experimental data, they are the most natural candidates for our need for specific u, π functions to evaluate (3), and yield the best answer we can hope to have to our question—without conducting a measurement of the functions u and π in the specific marketplace of interest—for ‘real’ population preferences u, π .

Cumulative prospect theory preferences u, π . We will first investigate the condition in Theorem 4.1 for the u, π functions from Tversky and Kahneman [1992], which uses experimental data to estimate the value functions $u(\cdot)$ and the decision weights $\pi(\cdot)$.

The value function u over payoffs proposed by Tversky and Kahneman [1992] is

$$u(x) = \begin{cases} x^\alpha & x \geq 0 \\ -\lambda(-x)^\alpha & x < 0. \end{cases} \quad (4)$$

The exponent $\alpha \leq 1$ models risk aversion since x^α is strictly concave for $\alpha \in (0, 1)$, with smaller α indicating a greater degree of risk aversion.¹⁹ The multiplier $\lambda > 1$ models loss aversion—a loss of x incurs greater disutility $|u(-x)| = \lambda x^\alpha$ than the corresponding utility $u(x) = x^\alpha$ from a gain of x —with a larger value of λ indicates a greater degree of loss aversion.

The decision weight function π proposed in Tversky and Kahneman [1992] is

$$\pi(p) = \frac{p^\delta}{[p^\delta + (1-p)^\delta]^{1/\delta}}, \quad (5)$$

where the parameter $\delta \in [0, 1]$ models the extent to which agents use decision weights that differ from outcomes’ probabilities. As δ becomes smaller, there is increased deviation from probability weighting, while at $\delta = 1$, this function reduces to $\pi(p) = p$, corresponding to expected utility maximizing preferences with decision weights equal to the probabilities of outcomes.

We can now use Theorem 4.1 with these functions to ask whether and when a contest dominates a fixed-payment contract for a population with preferences describes by these payoffs and decision weights u, π . Our main result in this section, Theorem 4.4 below, shows that if the parameter α describing the degree of risk aversion in u is larger than the parameter δ describing the degree of probability weighting in π , a contest will eventually dominate a fixed payment scheme for large enough total payouts W . Further, the total payment $W^*(\lambda, \alpha, \delta)$ beyond which the contest dominates behaves as one might expect with the parameters quantifying risk and loss aversion: the minimum contest prize W^* needed to beat a fixed-payment contract increases as α becomes smaller—corresponding to more risk-averse agents—and as λ becomes larger, which corresponds to more loss aversion.

¹⁹For this utility function the measure of relative risk aversion is $1 - \alpha$.

THEOREM 4.4. *Consider a population of agents with value functions and decision weights (u, π) as in (4) and (5). A contest dominates a fixed-payment contract for such a population for all large enough W if the parameters δ and α satisfy $\alpha > \delta$.*

In particular, if $\alpha > \delta$, the winner-take-all contest with prize W dominates a fixed-payment contract with total payment W for all $W \geq W^(\lambda, \alpha, \delta) = c(\lambda \cdot 2^{\frac{1-\delta}{\delta}})^{\frac{1}{\alpha-\delta}}$.*

(Note that while the theorem is stated in terms of W , it is really the *relative* scale W/c —the size of the prize relative to the cost that must be sunk to attempt to gain it—that determines when contests begin to dominate fixed payments. So a principal willing to offer a large sum W in absolute dollar amounts might still be better off with a fixed payment contract than a contest if working on the task for his contest involves a high cost c , and conversely for low absolute payouts W and ‘easy’ tasks with low c .)

The condition $\alpha > \delta$ in Theorem 4.4 has an intuitive interpretation as capturing the tradeoff between the increased weighting of small probabilities by π (which, roughly speaking, favors a contest) and agents’ risk and loss aversion (favoring the ‘safer’ fixed-payment contract). Recall that risk aversion increases as α shrinks further away from 1 (with $\alpha = 1$ corresponding to risk-neutrality), and deviation from linear probability weighting also increases as δ shrinks further away from 1 (with $\delta = 1$ corresponding to expected-utility behavior). The theorem says that if the degree of deviation from linear probability weighting in π is ‘greater’ than the extent of deviation from risk-neutrality in u as captured by the condition $\alpha > \delta$, a contest will eventually dominate fixed payment contracts for the agent population with preferences u, π .

We note that the values of α and δ estimated in Tversky and Kahneman [1992] are $\alpha = 0.88$ and $\delta = 0.65$, so that $\alpha - \delta$ is indeed positive according to these estimated cumulative prospect theory valuations. Theorem 4.4 therefore suggests that contests will indeed serve as a better incentive scheme than fixed-payment contracts for large enough W .

Alternate functional forms. Subsequent work on econometric estimation of preferences in risky decision-making environments corroborates the main findings of cumulative prospect theory, although using somewhat different functional forms for u and π than those in (4) and (5) from Tversky and Kahneman [1992]. For completeness, we outline how our qualitative conclusions in Theorem 4.4 continue to hold for these alternate functional forms as well. We consider the model from Bruhin et al. [2010], which performs an econometric estimation based on experimental results from an extensive study with 18000 subjects. The π and u functions in Bruhin et al. [2010] are

$$\pi(p) = \frac{\eta p^\gamma}{\eta p^\gamma + (1-p)^\gamma}; \quad u(x) = \begin{cases} x^\alpha & x \geq 0 \\ -\lambda(-x)^\beta & x < 0. \end{cases} \quad (6)$$

where the parameter γ in the π function plays a role akin to δ in the model from Tversky and Kahneman [1992], with small values of $\gamma \in (0, 1)$ indicating greater deviation from linear weighting of probabilities, and the additional parameter β in u allows different curvatures for the utility from losses and gains. Again, we can use these functional forms in Theorem 4.1 to address when contests dominate a fixed-payment contract, giving us the analog of Theorem 4.4:

THEOREM 4.5. *Consider a population of agents with value functions and decision weights (u, π) as in (6). A contest dominates a fixed-payment contract for such a population for all large enough W if the parameters γ and α satisfy $\alpha > \gamma$.*

Bruhin et al. [2010] find that the majority of their experimental subjects (~80%) display significant non-linear weighting of probabilities, segregating consistently into two subpopulations—CPT-I and CPT-II—with lesser and greater deviation from EU preferences; they estimate $\alpha = 0.901$ and $\gamma = 0.309$ for the CPT-I and $\alpha = 0.957$ and $\gamma = 0.451$ for the CPT-II population. Again, these parameter values satisfy the condition ($\alpha > \gamma$) under which a contest can dominate fixed-payment contracts in a market described by these preferences.

Theorem 4.4 and Theorem 4.5 allow us to affirmatively answer the question posed at the start of this section. Together with the estimated values of α and δ from Tversky and Kahneman [1992] and α and γ from Bruhin et al. [2010], these results suggest that—to the extent that these parameters indeed describe the decision-making behavior of agents in online crowdsourcing environments²⁰—a principal who values the output from crowdsourcing his task adequately highly (compared to the cost to a worker to produce that output) might indeed choose a different kind of mechanism, and do better as a result, if he accounts for deviations from the standard economic model of expected utilities that are typically used in theoretical design.

5. CONCLUSION AND FURTHER DIRECTIONS

In this paper, we asked whether behavioral biases—in particular, the consistently observed deviations of behavior from expected utility theory, the standard economic model of decision making under uncertainty—can matter to mechanism design. We explored this problem in the context of the incentive design problem faced by a principal in a crowdsourcing market with participation-only strategic choices, and demonstrated that a principal’s choice of contract can change in fundamental way—from an output-independent, ‘safe’ fixed-payment scheme to a riskier, output-contingent contest—depending on whether agents behave according to expected utility theory (with risk aversion) or prospect theory preferences. While our model is a reasonable description of the specific crowdsourcing markets we use to motivate it, it is deliberately chosen to be the simplest, most parsimonious possible model that illustrates our central point—that deviations from expected utility theory, according to prospect theory preferences, are potentially quite significant to theoretical design—in a realistic setting. This means that there are a number of complexities of crowdsourcing markets that our model does not capture, presenting promising, and important, directions for further work; we discuss several of these in the full version of the paper.

Prospect theory in mechanism design: Beyond principal-agent problems. The question of how choice according to prospect theory versus expected utility affects equilibrium analysis and optimal design extends to domains beyond principal-agent problems in online labor—principal-agent problems are merely one amongst many applications of mechanism design where agents make decisions under uncertainty. A number of other domains to which mechanism design applies, especially those motivated by online systems where individuals rather than large firms are decision makers—such as online auctions with small bidders, or reputation systems in peer-to-peer economies—depend on the decisions of individual agents to whom these behavioral biases do apply. The issue of whether agents choose according to prospect theory versus expected utility preferences is a fundamental component of equilibrium analysis and the optimal design question in each of these environments, and as such, lays open a vast field of problems for further exploration.

REFERENCES

- Adrian Bruhin, Helga Fehr-Duda, and Thomas Epper. 2010. Risk and Rationality: Uncovering Heterogeneity in Probability Distortion. *Econometrica* 78, 4 (2010).
- H. Chang, J. Cvitanic, and X. Zhou. 2014. Optimal Contracting with Moral Hazard and Behavioral Preferences. *Working paper* (2014).
- Yeon-Koo Che and Ian Gale. 2003. Optimal Design of Research Tournaments. *American Economic Review* 93, 3 (2003).
- David de Meza and David C. Webb. 2007. Incentive Design under Loss Aversion. *Journal of the European Economic Association* 5, 1 (2007).
- Richard L. Fullerton, Bruce G. Linstler, Michael McKee, and Stephen Slate. 2002. Using Auctions To Reward Tournament Winners: Theory and Experimental Investigations. *RAND Journal of*

²⁰It is reasonable to imagine that the values of α and δ for a population could well be sensitive to the (scale of the) payoffs offered, the population of agents and their reference wealth levels, and so on, as well as the setup yielding the measurement.

- Economics* 33, 1 (Spring 2002).
- A. Ghosh and P. Hummel. 2015. Cardinal Contests. In *24th International World Wide Web Conference (WWW'15)*.
- Arpita Ghosh and Robert Kleinberg. 2014. Behavioral Mechanism Design: Optimal Contests for Simple Agents. In *Proc. 15th ACM Conference on Economics and Computation (EC)*. 913–930.
- Jerry R. Green and Nancy L. Stokey. 1983. A Comparison of Tournaments and Contracts. *Journal of Political Economy* 91, 3 (1983), 349–364.
- Oliver Hart and Bengt Holmstrom. 1987. The Theory of Contracts. In *Advances in Economic Theory, Fifth World Congress*.
- Grace YoungJoo Jeon, Yong-Mi Kim, and Yan Chen. 2010. Re-examining Price As a Predictor of Answer Quality in an Online Q&A Site. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*.
- Daniel Kahneman and Amos Tversky. 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica* 47, 2 (1979).
- Ravi Kanbur, Jukka Pirttil, and Matti Tuomala. 2008. Moral Hazard, Income Taxation and Prospect Theory*. *Scandinavian Journal of Economics* 110, 2 (2008).
- Robert Kraut and Paul Resnick. 2012. *Building successful online communities: Evidence-based social design*. Cambridge University Press.
- Edward P. Lazear and Sherwin Rosen. 1981. Rank-Order Tournaments as Optimum Labor Contracts. *Journal of Political Economy* 89, 5 (1981), 841–864.
- W. Mason and D. Watts. 2009. Financial Incentives and the "Performance of Crowds". In *Proceedings of the ACM SIGKDD Workshop on Human Computation (HCOMP '09)*.
- Daisuke Nakajima. 2011. First-price auctions, Dutch auctions, and buy-it-now-prices with Allais paradox bidders. 6, 3 (2011).
- Leonard Savage. 1954. *The Foundations of Statistics*. Wiley.
- Anja Schottner. 2008. Fixed-prize tournaments versus first-price auctions in innovation contests. 35, 1 (2008).
- Amos Tversky and Daniel Kahneman. 1992. Advances in Prospect Theory: Cumulative Representation of Uncertainty. *Journal of Risk and Uncertainty* 5 (1992).
- John von Neumann and Oscar Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton University Press.